Exhibit C

AskF5

K55502976: BIG-IP LTM-DNS operations guide | Chapter 6: BIG-IP DNS/DNS services

Operations Guide

Original Publication Date: Oct 09, 2018
Updated Date: Mar 03, 2020

Table of contents | << Previous chapter | Next chapter >>

This document reviews BIG-IP DNS offerings available from F5.

Contents

Chapter sections

- DNS Services features
- Upgrading to BIG-IP DNS 12.0 and later
    - Prerequisites
- BIG-IP DNS/DNS services basics
- BIG-IP DNS/DNS services core concepts
    - Configuration synchronization
    - BIG-IP DNS listeners
    - Data centers and virtual servers
    - Links
    - DNS Express
    - DNS Anycast
    - DNS cache
    - Hardware acceleration
    - DNSSEC
    - Auto-discovery
    - Address translation
    - ZoneRunner
    - iRules
    - iControl and iControl REST
- BIG-IP DNS load balancing

    - Monitors
    - Wide IPs
    - Record types
    - Wide IPs and pools

- BIG-IP DNS Minimal Response Setting
    - Examples
    - Effects on DNS Performance
    - Load balancing logic
    - Topology

- BIG-IP DNS Architectures
    - Delegated mode
    - Screening mode
    - Replacing a DNS server

- BIG-IP DNS iQuery
    - iQuery Agents
    - big3d software versioning
    - Determining the health of the iQuery mesh

- BIG-IP DNS device service clustering
- BIG-IP DNS query logging
- BIG-IP DNS Statistics

Figures

- Figure 6.1: Cache miss
- Figure 6.2: Cache hit
- Figure 6.3: Resolver DNS cache
- Figure 6.4: DNSSEC chain of trust
- Figure 6.5: wide IP with A and AAAA virtual servers in a single pool
- Figure 6.6: wide IP with CNAME pool configuration
- Figure 6.7: “A” Resource Record configuration
- Figure 6.8: “AAAA” Resource Record configuration
- Figure 6.9: CNAME For Type A Resource Record
- Figure 6.10: CNAME For Type A Resource Record with Static Target
- Figure 6.11: Type MX wide IP
- Figure 6.12: SRV wide IPs configuration
- Figure 6.13: NAPTR type wide IP configuration
- Figure 6.14: BIG-IP LTM/BIG-IP DNS configuration

DNS Services features

BIG-IP DNS (formerly BIG-IP GTM) is a DNS-based module which monitor the availability and performance of global resources, such as distributed applications, in order to control network traffic patterns.

DNS caching is a DNS services feature that can provide responses for frequently requested DNS records from a cache maintained in memory on BIG-IP systems. This feature can be used to replace or reduce load on other DNS servers.

DNS Express is a DNS Services feature that allows the BIG-IP system to act as an authoritative slave server. DNS Express relies on NOTIFY, AXFR, and IXFR to transfer zone data from a master authoritative server and store it in memory for high performance. DNS Express does not offer DNS record management capability.

Response Policy Zones (RPZ) allow the BIG-IP system, when configured with a DNS cache, to filter DNS requests based on the resource name being queried. This can be used to prevent clients from accessing known-malicious sites.

Upgrading to BIG-IP DNS 12.0 and later

Prerequisites

F5 requires BIG-IP 11.x to be installed before upgrading. Upgrades to BIG-IP DNS 12.0 and later from versions prior to 11.x are not supported.

TMOS Shell (**tmsh**) scripts on BIG-IP 11.x and earlier use DNS pool and wide IP commands do not work correctly. They must be converted to use the new pool and wide IP Type parameters.

If any DNS objects are currently active in the configuration DNS, configuration files created on BIG-IP devices prior to BIG-IP 12.0 (for example bigip_DNS.conf, UCS, SCF) cannot be loaded on a BIG-IP device running BIG-IP 12.0.

To load the DNS configuration, the DNS configuration must be blank. Load an empty **bigip_DNS.conf** and save the configuration. After the blank configuration loads, a **bigip_DNS.conf**, UCS, or SCF created in a prior version can be loaded.

It is possible to use the deprecated DNS iControl API interfaces pool, wide IP, application, and pool member commands on a device running BIG-IP 12.0.

If the configuration has not been modified to use any of the newly supported types, F5 recommends that you use the DNS iControl API interfaces Poolv2, wide IPv2.

BIG-IP DNS/DNS services basics

BIG-IP DNS is the module built to monitor the availability and performance of global resources and use that information to manage network traffic patterns.

You can use BIG-IP DNS module to do the following:

- Direct clients to local servers for globally-distributed sites using a GeoIP database.
- Change the load balancing configuration according to current traffic patterns or time of day.
- Set up global load balancing among disparate BIG-IP LTM systems and other hosts.
- Monitor real-time network conditions.
- Integrate a content delivery network from a CDN provider.

To implement BIG-IP DNS, you need to understand the following terminology and basic functionality:

- Configuration synchronization (ConfigSync) ensures the rapid distribution of BIG-IP DNS settings among BIG-IP DNS systems in a synchronization group.

- Load balancing divides work among resources so that more work gets done in the same amount of time and, in general, all users get served faster. BIG-IP DNS selects the best available resource using either a static or a dynamic load balancing method. When using a static load balancing method, BIG-IP DNS selects a resource based on a pre-defined pattern. When using a dynamic load balancing method, BIG-IP DNS selects a resource based on current performance metrics.
- Prober pool is an ordered collection of one or more BIG-IP systems that can be used to monitor specific resources.
- wide IP is a mapping of a fully-qualified domain name (FQDN) to a set of virtual servers that host the domains content, such as a web site, an e-commerce site, or a content delivery network (CDN). BIG-IP DNS intercepts requests for domain names that are wide IPs and answers them based on the wide IP configuration.
- iQuery is an XML protocol used by BIG-IP DNS to communicate with other BIG-IP systems.
- BIG-IP DNS listener is a specialized virtual server that provides DNS services.
- Probe is an action the BIG-IP system takes to acquire data from other network resources. BIG-IP DNS uses probes to track the health and availability of network resources
- Data center is where BIG-IP DNS consolidates all the paths and metrics data collected from the servers, virtual servers, and links.
- Virtual server is a combination of IP address and port number that points to a resource that provides access to an application or data source on the network.
- Link is a logical representation of a physical device (router) that connects the network to the Internet.
- Domain Name System Security Extensions (DNSSEC) is an industry-standard protocol that functions to provide integrity for DNS data.

For more information, refer to BIG-IP DNS load balancing.

BIG-IP DNS/DNS services core concepts

This section covers DNS Express, DNS cache, auto-Discovery, address translation, and ZoneRunner.

Configuration synchronization

Configuration synchronization (ConfigSync) ensures the rapid distribution of BIG-IP DNS settings to other BIG-IP DNS systems that belong to the same synchronization group. A BIG-IP DNS synchronization group might contain both BIG-IP DNS and BIG-IP Link Controller systems.

ConfigSync occurs in the following manner:

- When a change is made to a BIG-IP DNS configuration, the system broadcasts the change to the other systems in BIG-IP DNS synchronization group.
- When a configuration synchronization is in progress, the process must either complete or time out before another configuration synchronization can occur.

***Note**: It is important to have a working Network Time Protocol (NTP) configuration because BIG-IP DNS relies on timestamps for proper synchronization.*

BIG-IP DNS listeners

A listener is a specialized virtual server that provides DNS services on port 53 and at the IP address assigned to the listener. When a DNS query is sent to the listener, BIG-IP DNS either handles the request locally or forwards the request to the appropriate resource.

BIG-IP DNS responds to DNS queries on a per-listener basis. The number of listeners created depends on the network configuration and the destinations to which specific queries are to be sent. For example, a single BIG-IP DNS can be the primary authoritative server for one domain, while forwarding other DNS queries to a different DNS server. BIG-IP DNS always manages and responds to DNS queries for the wide IPs that are configured on the system.

Data centers and virtual servers

All of the resources on a network are associated with a data center. BIG-IP DNS consolidates the paths and metrics data collected from the servers, virtual servers, and links in the data center. BIG-IP DNS uses that data to conduct load balancing and route client requests to the best-performing resource, based on a variety of factors.

BIG-IP DNS may send all requests to one data center when another data center is down. This may work well for disaster recovery sites.

Alternatively, BIG-IP DNS might send a request to the data center that has the fastest response time.

Or, BIG-IP DNS may send a request to the data center that is located closest to the client's source address. For example, the system may send a client located in France to a host also located in France rather than the United States, greatly reducing traffic round-trip times.

***Note**: The resources associated with a data center are available only when the data center is also available.*

Virtual servers

A virtual server is a specific IP address and port number that points to a resource on the network. In the case of host servers, this IP address and port number likely point to the resource itself. With load balancing systems, virtual servers are often proxies that allow the load balancing server to manage a resource request across a large number of resources.

***Notes**:*

- *You can configure virtual server status to be dependent only on the timeout value of the monitor associated it. In a multi-bladed environment, this configuration ensures that when the primary blade in a cluster becomes unavailable, the **gtmd** agent on the new primary blade has time to establish new iQuery connections with and receive updated status from other BIG-IP systems.*
- *The **big3d** agent on the new primary blade has 90 seconds to run (the timeout value of the BIG-IP monitor) before it times out.*

Links

A link is an optional BIG-IP DNS or BIG-IP Link Controller configuration object which represents a physical device that connects a network to the Internet. BIG-IP DNS tracks the performance of links. Performance results influence the availability of pools, data centers, wide IPs, and distributed applications.

When you create one or more links, the BIG-IP system uses the following logic to automatically associate virtual servers with the link objects:

- BIG-IP DNS and BIG-IP Link Controller associate the virtual server with the link by matching the subnet addresses of the virtual server, link, and self IP address. Most of the time, the virtual server is associated with the link that is on the same subnet as the self IP address.
- In some cases, BIG-IP DNS and BIG-IP Link Controller cannot associate the virtual server and link because the subnet addresses do not match. When this occurs, the system associates the virtual server with the default link which is assigned to the data center. This association may cause issues if the link that is associated with the virtual server does not provide network connectivity to the virtual server.

- If the virtual server is associated with a link that does not provide network connectivity to that virtual server, BIG-IP DNS and BIG-IP Link Controller may incorrectly return the virtual server IP address in the DNS response to a wide IP query even if the link is disabled or marked as down.

DNS Express

DNS Express enables the BIG-IP system to function as a replica authoritative nameserver and answer DNS queries at high speeds. However, since DNS Express doesn't use BIND DNS software, it doesn't have the same security vulnerabilities as a typical BIND implementation.

DNS Express supports the standard DNS NOTIFY protocol from primary authoritative nameservers and uses the AXFR/IXFR mechanism to transfer zone data. The primary authoritative nameserver is not listed in the start of authority (SOA) of the zone data, and is therefore protected, or hidden.

Optionally, transaction signature (TSIG) may be used to secure the zone data transferred from the primary nameserver.

DNS Express doesn't support modifying records. Instead, records are modified on the primary nameserver and DNS Express is notified of the changes. However, the BIG-IP system may be configured as the primary authoritative nameserver using Zonerunner.

DNS Anycast

You can configure IP Anycast for DNS services on BIG-IP systems that have the advanced routing module license. Anycast describes a one-to-nearest communication between a client and the nearest recipient within a group. The routing protocol directs client queries to a recipient in the target group based on the routing algorithm for the specified protocol. This capability improves reliability and performance, while distributing load across multiple BIG-IP systems.

To enable IP Anycast for DNS services, ZebOS dynamic routing needs to be enabled and configured with the appropriate protocol for your deployment. Then, listeners must be configured on each of the BIG-IP systems with the shared IP Anycast address and route advertisement enabled under the advanced settings of the listener.

DNS Anycast benefits

DNS queries are sent to the Anycast IP address that is defined on multiple BIG-IP systems. If a system becomes unavailable, the route to that system is removed dynamically.

Performance is improved by routing queries to the nearest BIG-IP system.

Distributing the load across multiple, geographically distributed BIG-IP systems helps mitigate distributed denial-of-service attacks (DDoS).

DNS cache

The DNS cache feature is available as a DNS add-on module for BIG-IP LTM. DNS Cache has three different configurable forms of DNS cache: transparent, resolver, and validating resolver.

Transparent DNS cache

The transparent cache object is configurable on the BIG-IP system to use external DNS resolvers to resolve queries, and then cache the responses from the multiple external resolvers. When a consolidated cache is in front of external resolvers (each with their own cache), it can produce a much higher cache hit percentage.

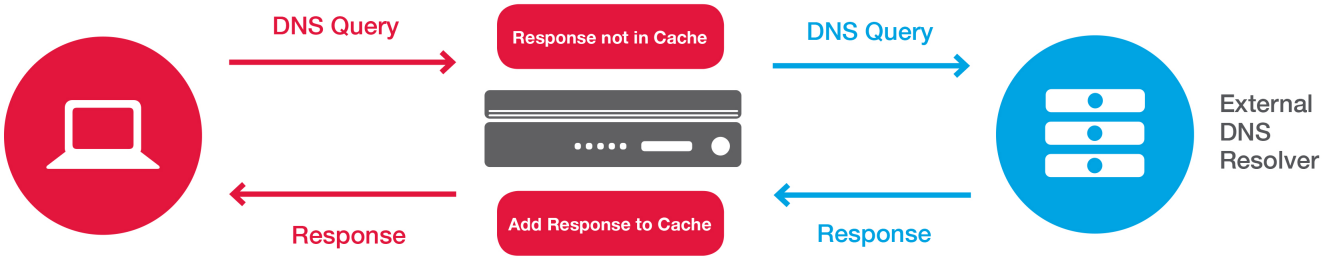

Figure 6.1 Cache miss

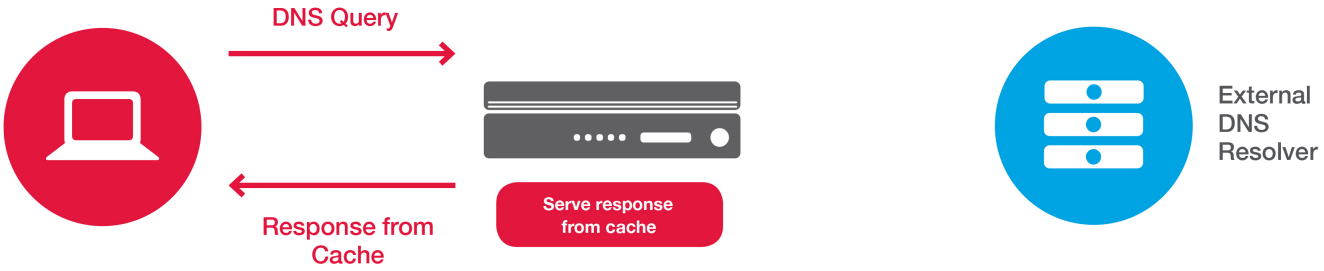

Figure 6.2: Cache hit

***Note**: The transparent cache contains messages and resource records.*

F5 recommends that you configure the BIG-IP system to forward queries which cannot be answered from the cache to a pool of local DNS servers, rather than to the local BIND instance because BIND performance is slower than using multiple external resolvers.

***Note**: For systems using the DNS Express feature, if DNS Express is available to answer, no further processing is needed.*

Resolver DNS Cache

You may configure a resolver cache on the BIG-IP system to resolve DNS queries and cache the responses. The next time the system receives a query for a response that exists in the cache, the system returns the response from the cache. The resolver cache contains messages, resource records, and the nameservers the system queries to resolve DNS queries.

***Note**: It is possible to configure the local BIND instance on the BIG-IP system to act as an external DNS resolver. However, the performance of BIND is slower than using a resolver cache.*

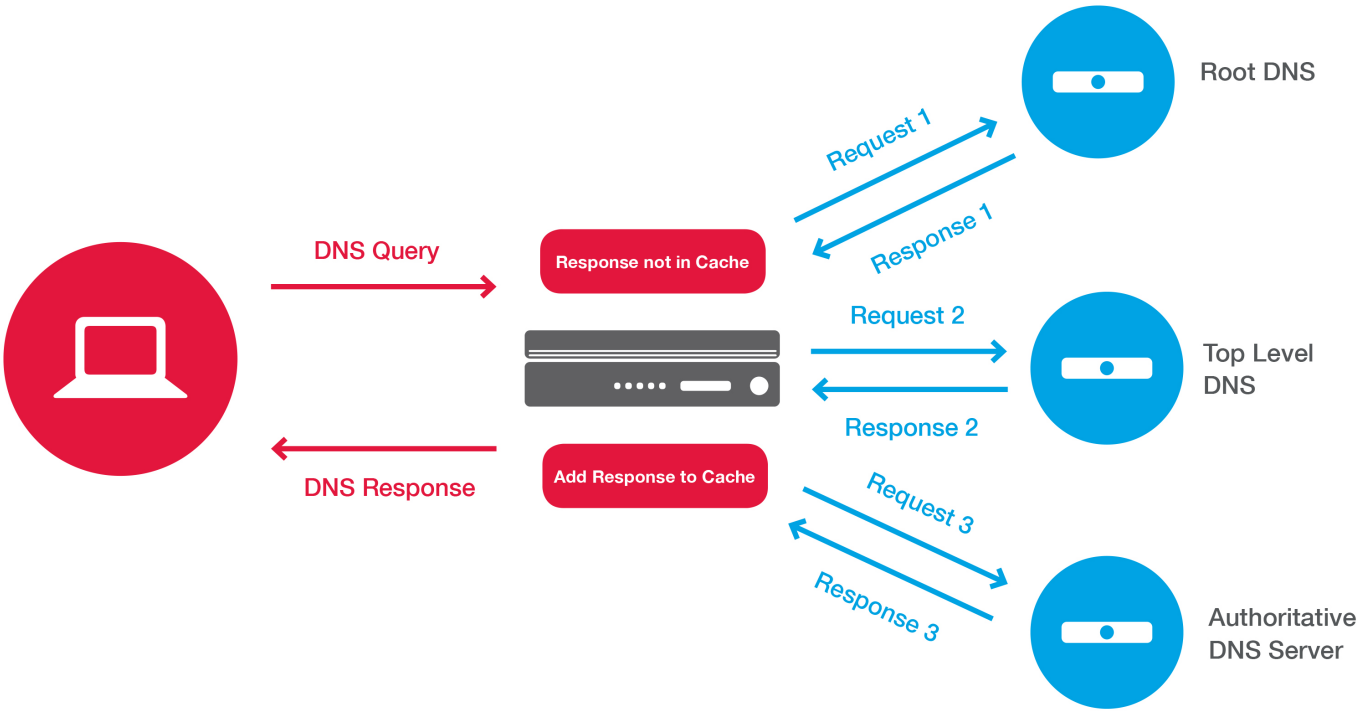

Figure 6.3 Resolver DNS cache

Validating resolver DNS cache

The validating resolver DNS cache may be configured to recursively query public DNS servers, validate the identity of the DNS server sending the responses, and then cache the responses. The next time the system receives a query for a response that exists in the cache, the system returns the DNSSEC-compliant response from the cache.

The validating resolver cache contains messages, resource records, the nameservers the system queries to resolve DNS queries, and DNSSEC keys.

For more information about setting up each of the DNS express caching methodologies, refer to ***DNS Cache: Implementations***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

DNS cache optimization

DNS Cache optimization can be complex. You want to maximize the number of DNS cache hits while conserving allocated memory.

There are diminishing returns to cache too much. An average cache has a hit rate of 80-85 percent, and a majority of those records typically require less than 1 gigabyte (GB) of space. If your cache is larger than 1 GB, you may be taking up memory space by caching objects that are not requested often.

To optimize your cache, set it for a few hundred MB, observe the cache, and then make adjustments as necessary.

Viewing and adjusting cache size using the Configuration utility

1. Go to **DNS** > **Caches** > **Cache List**.
2. Select the name of the cache.
3. Note the **Message Cache** size.

***Note**: This setting is in bytes and is a per Traffic Management Microkernel (TMM) setting.*

4. Multiply the number of CPU cores by this value to get the total memory allocation space.
5. Note the **Resource Record Cache** size.
   ***Note**: This setting is in bytes and is a per TMM setting.*
6. Multiply the number of CPU cores by this value to get the total memory allocation space..
7. Set the **Nameserver Cache Count** value as needed.
8. In the **Unsolicited Reply Threshold** box, change the default value if you are using the BIG-IP system to monitor for unsolicited replies using SNMP.
   ***Note**: The BIG-IP system always rejects unsolicited replies.*

   The default value of 0 (off) indicates the system does not generate SNMP traps or log messages when rejecting unsolicited replies.

   Changing the default value alerts you to a potential security attack, such as cache poisoning or DDoS. For example, if you specify 1,000,000 unsolicited replies, each time the system receives 1,000,000 unsolicited replies, it generates an SNMP trap and log message.

   Check the number of cache hits that you receive as seen in the cache statistics. Clear the counters and check again during peak usage hours.

   Observe and record the percentage of cache hits in relation to the total cache size.
9. Decrease the cache settings and check the cache statistics again. Note whether or not the number of cache hits remain the same.
10. Continue to adjust the cache settings until you reach an optimal balance between the numbers of cache hits versus cache sizing.
    iRules may also be used to manipulate a DNS response given from cached resource records. For example:

    Assume that we wish to decrease the number of repeat DNS cache hits if the number of available pool members for a Transparent Cache drops beneath a given number. This goal can be achieved if we double the TTL given to a client if the number of pool members drops beneath x.

Hardware acceleration

For BIG-IP LTM 12.0 running the appropriate hardware, DNS response cache and protocol validation can be accelerated in hardware to lighten load, help mitigated DDoS attacks, and improve response time for DNS queries.

Disabled by default, these option scan be enabled on the DNS Profile of the UDP listener with the following prerequisites:

- BIG-IP LTM 12.0 or later installed on the appropriate hardware.
- Intelligent L7 Bitstream is enabled on the BIG-IP system.

***Note**: Hardware acceleration features only work on BIG-IP LTM 12.x used with B2250 blades in the 2x00 series VIPRION chassis with L7 Intelligent Bitstream enabled. The BIG-IP system does not prevent you from enabling the features, even if your system does not meet these criteria. Enabling them without the appropriate criteria has no effect on your BIG-IP system.*

DNS response cache in hardware

The DNS Response cache is designed to service millions of queries per second. It caches responses from any of the software caches listed in the previous section as well as DNS Express. Given their dynamic nature, GTM responses are never cached. The Hardware cache is flushed every second, requiring another DNS query.

Protocol validation in hardware

By moving some protocol validation into hardware, system performance is enhanced by quickly discarding malformed DNS queries. This relieves the burden on the CPU and provides increased DDoS protection.

Troubleshooting

If these options do not function, make sure of the following:

- You have BIG-IP 12.0 or later installed on the appropriate hardware.
- You have enabled L7 Intelligent Bitstream.
- You have enabled the options in the DNS Profile.

***Note**: To enable these options on your DNS profile, F5 recommends making a copy of the default profile and updating the copy, leaving the default unchanged.*

With the Hardware Acceleration features enabled, most queries are answered from cache.

Viewing statistics in the Configuration utility

- Go to **Statistics** > **Module Statistics** > **DNS** > **Delivery**.

DNSSEC

DNSSEC is an extension to the Domain Name Service (DNS) that ensures the integrity of data returned by domain name lookups by incorporating a chain of trust in the DNS hierarchy. DNSSEC provides origin authenticity, data integrity and secure denial of existence.

Specifically, origin authenticity ensures that resolvers can verify that data has originated from the correct authoritative source. Data Integrity verifies that responses are not modified in-flight, and Secure Denial of Existence ensures that when there is no data for a query, that the authoritative server can provide a response that proves no data exists.

The basis of DNSSEC is public key cryptography (PKI). A chain of trust is built with public-private keys at each layer of the DNS architecture.

DNSSEC key types

DNSSEC uses two kinds of keys: key-signing keys and zone signing keys.

- Key signing key is used to sign other keys in order to build the chain of trust. This key is sometimes cryptographically stronger and has a longer lifespan than a Zone signing key.
- Zone signing key is used to sign the data that is published in a zone. DNSSEC uses the key signing keys and zone signing keys to sign and verify records within DNS.

DNSSEC chain of trust

When a user requests a site, DNS translates the domain name into an IP address through a series of recursive lookups that form a “chain” of requests. Each stop in this chain inherently trusts the other parts of the chain, and this trust may be exploited by an attack.

For example, if an attack manipulates a servers or some traffic along the chain, it can redirect the client to a website where malware is waiting.

DNSSEC mitigates this problem by validating the response of each part of the chain with digital signatures. These signatures help build a “chain of trust” that DNS can rely on when answering requests. To form the chain of trust, DNSSEC starts with a “trust anchor” and everything following that trust anchor is trusted. Ideally, the trust anchor

is the root zone.

ICANN published the root zone trust anchor, and root operators began serving the signed root zone in July 2010. With the root zone signed, all other zones following it can also be signed, thus forming a solid and complete chain of trust. Additionally, ICANN also lists the Top Level Domains that are currently signed and have trust anchors published as DS records in the root zone.

The following figure shows the building blocks for the chain of trust from the root zone.

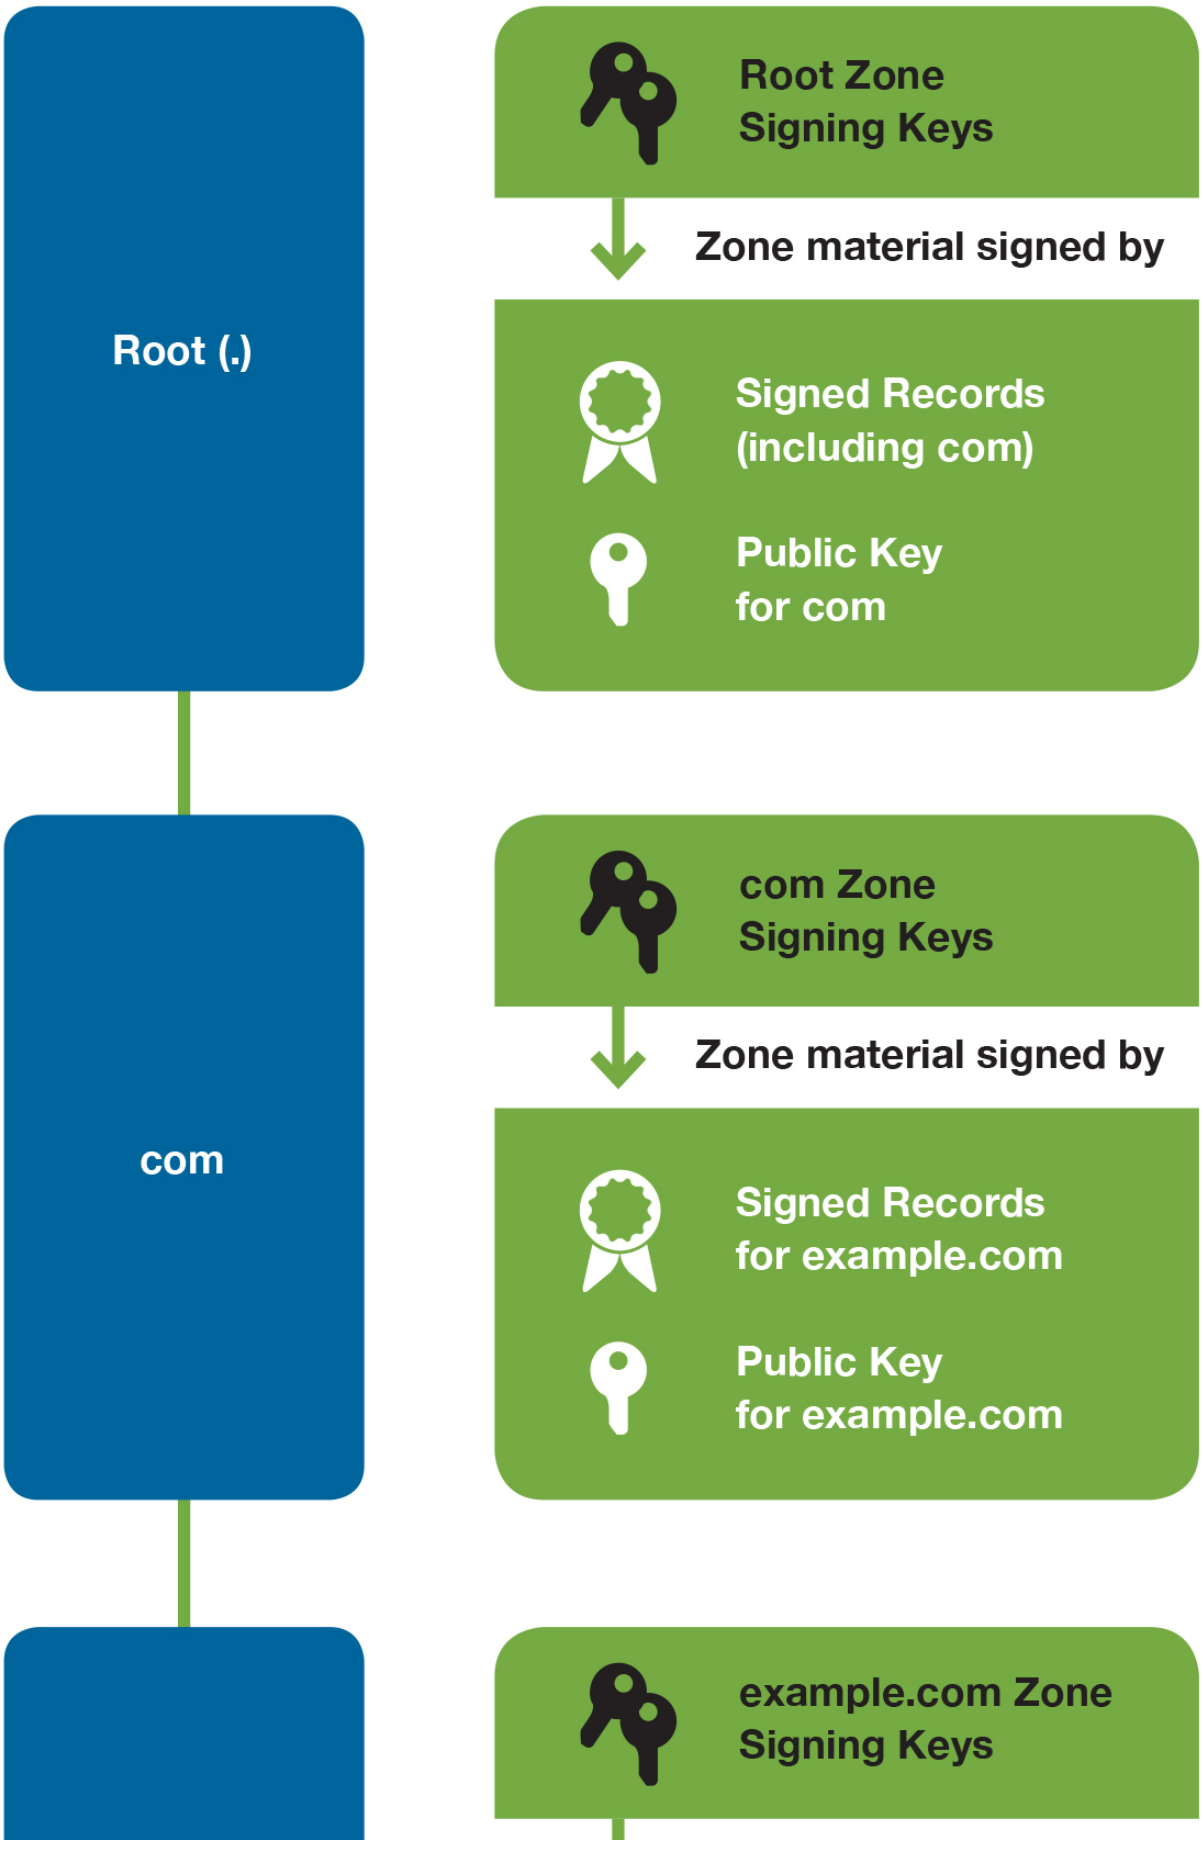
Root (.)
Root Zone Signing Keys
Zone material signed by
Signed Records (including com)
Public Key for com
com
com Zone Signing Keys
Zone material signed by
Signed Records for example.com
Public Key for example.com
example.com Zone Signing Keys

Figure 6.4 DNSSEC chain of trust

For more information, refer to **Configuring DNSSEC** in ***BIG-IP DNS Services***.

Auto-discovery

Auto-discovery is a process through which BIG-IP DNS automatically identifies resources that it manages. BIG-IP DNS can discover two types of resources: virtual servers and links.

Each resource is discovered on a per-server basis, so you can employ auto-discovery only on the servers you specify.

The auto-discovery feature of BIG-IP DNS has three modes that control how the system identifies resources:

- Disabled: BIG-IP DNS does not attempt to discover any resources. Auto-discovery is disabled on BIG-IP DNS by default.
- Enabled: BIG-IP DNS regularly checks the server to discover any new resources. If a previously discovered resource cannot be found, BIG-IP DNS deletes it from the system.
- Enabled (No Delete): BIG-IP DNS regularly checks the server to discover any new resources. Unlike the Enabled mode, the Enabled (No Delete) mode does not delete resources, even if the system cannot currently verify their presence.

***Note**: Enabled and Enabled (No Delete) modes query the servers for new resources every 30 seconds by default.*

Important: Auto-discovery must be globally enabled at the server and link levels, and the frequency at which the system queries for new resources must be configured.

For information about enabling auto-discovery on virtual servers and links, refer to **Discovering resources automatically** in the ***Configuration Guide for BIG-IP Global Traffic Manager***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Address translation

Several objects in BIG-IP DNS allow the specification of address translation. Address translation is used in cases where the object is behind a Network Address Translation (NAT). For example, a virtual server may be known by one address on the Internet but another address behind the firewall. When configuring these objects, the address is the external address and is returned in any DNS responses generated by BIG-IP DNS.

When probing, the BIG-IP system may use either the address or translation, depending on the situation. As a general rule, if both the BIG-IP system performing the probe and the target of the probe are in the same data center and both have a translation, the probe uses the translations. Otherwise, the probe uses the address.

Specifying a translation on a BIG-IP server causes virtual server auto-discovery to silently stop working. This is because BIG-IP DNS has no way of knowing what the externally visible address should be for the discovered virtual server address, which is a translation. For more information, refer to K9138: BIG-IP GTM system disables virtual server auto-discovery for BIG-IP systems that use translated virtual server addresses.

ZoneRunner

ZoneRunner is an F5 product used for zone file management on BIG-IP DNS. You may use the ZoneRunner utility to create and manage DNS zone files and configure the BIND instance on BIG-IP DNS. With the ZoneRunner utility, your system can:

- Import and transfer DNS zone files.
- Manage zone resource records.
- Manage views.
- Manage a local nameserver and the associated configuration file, named.conf.
- Transfer zone files to a nameserver.
- Import only primary zone files from a nameserver.

BIG-IP DNS ZoneRunner utility uses dynamic update to make zone changes. All changes made to a zone using dynamic update are written to the zone's journal file.

Important: F5 recommends that you let the ZoneRunner utility manage the DNS/BIND file rather than manually editing the file. If manual editing is required, the zone files must be frozen to avoid issues with name resolution and dynamic updates.

To prevent the journal files from being synchronized if BIG-IP DNS is configured to synchronize DNS zone files, the zone must be frozen on all BIG-IP DNS systems.

For more information refer to **ZoneRunner** in ***BIG-IP GTM: Concepts***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

iRules

iRules can be attached to or associated with a wide IP, and in BIG-IP 11.5.0 and later it can be attached to or associated with the DNS listener.

When using iRules with BIG-IP DNS, there are two possible places to attach iRules: either to the wide IP or to the DNS listener. The iRules commands and events available depend on where the iRules are attached in the configuration. Some iRules functions require BIG-IP LTM to be provisioned alongside BIG-IP DNS.

BIG-IP DNS 12.0 and later includes changes to DNS iRules. Because a name is no longer sufficient to identify a wide IP, the iRules command syntax is updated to allow explicit declaration of the wide IP name and type. However, a wide IP type is not always required for iRules dealing with DNS. If a wide IP type is not specified, the type is automatically set during run-time to the type of the wide IP using iRules. Therefore, it is possible to have the same iRule attached to multiple wide IPs of different types and allow TMM to set the type as the iRule is executed.

tmsh

In BIG-IP 12.0 and later, **tmsh** commands support a new type attribute. The following shows the updates to command syntax, which includes the type parameters.

- `tmsh show DNS wide IP <A|AAAA|CNAME|MX|SRV|NAPTR> [wide IP Name]`
- `tmsh show DNS pool <A|AAAA|CNAME|MX|SRV|NAPTR> [Pool Name]`

iControl and iControl REST

The iControl API allows administrators to interact with their BIG-IP devices programmatically to update configurations, view statistics, etc. Two versions of the iControl API exist: iControl and iControl REST.

iControl changes in BIG-IP DNS 12.0 and later

BIG-IP 12.0 includes changes to both iControl and iControl REST. Both versions of iControl can interface with the DNS module on BIG-IP. The addition of the Type field when defining wide IPs necessitate a change to both iControl APIs. The support for existing customer iControl applications is limited and is explained in greater detail in the next two sections.

iControl

BIG-IP 12.0 and later supports iControl for DNS without any changes to existing applications. However, iControl only works with DNS configurations that do not use any of the new resource record types introduced in BIG-IP 12.0.

To learn more about iControl, visit F5 DevCentral iControl Wiki Home.

iControl REST

To learn more about iControl REST, visit the F5 DevCentral iControl REST Home.

BIG-IP 12.0 and later does not support iControl REST for DNS without changing the existing application. The iControl REST API functions are based on the underlying **tmsh** structure on the BIG-IP device. The **tmsh** commands for both wide IPs and pools change in this release; therefore, the iControl REST API changes.

BIG-IP DNS load balancing

Monitors

BIG-IP DNS uses health monitors to determine the availability of the virtual servers used in its responses to DNS requests. Detailed information about monitors can be found in ***BIG-IP Global Traffic Manager: Monitors Reference***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Probers

When running a monitor, BIG-IP DNS may request that another BIG-IP device probe the target of the monitor. BIG-IP DNS may choose a BIG-IP system in the same data center with the target of the monitor to actually send the probe and report back the status. This can minimize the amount of traffic that traverses the WAN.

The external and scripted monitors listed in the documentation use an external file to check the health of a remote system. BIG-IP DNS may request that another BIG-IP system in the configuration run the monitor. In order for the monitor to succeed, the remote BIG-IP system must have a copy of the external file.

Prober enhancements

BIG-IP DNS uses probers to decide if pool members are available based on algorithm criteria and then mark those resources up or down based on the results. In BIG-IP 13.0 and later, you can configure prober criteria to better control your resources. While in previous versions prober selection is performed automatically, in BIG-IP 13.0 and later, you can control the prober selection using **Prober Preference** and **Prober Fallback** options, in addition to the existing **Prober Pool** option.

BIG-IP 13.0 also enhances the **Availability Requirements** setting by adding a **Require** option which allows you to configure the number of probers using the same monitor at the same time. In earlier versions, if any prober fails then the resource is marked down but now you can configure the monitor probers so that the resource is marked down only if the configured number of probers are unsuccessful. These new monitor rules are available for BIG-IP DNS servers, virtual servers, pools, and pool members.

bigip monitor

The bigip monitor can be used to monitor BIG-IP systems. It uses the status of virtual servers determined by the remote BIG-IP system rather than sending individual monitor probes for each virtual server. It is recommended that BIG-IP LTM be configured to monitor its configuration elements so that it can determine the status of its virtual servers. This virtual server status is reported through the bigip monitor back to BIG-IP DNS. This is an efficient and effective way to monitor resources on other BIG-IP systems.

Application of monitors

In BIG-IP DNS configuration, monitors can be applied to the server, virtual server, pool and pool member objects. The monitor defined for the server is used to monitor all of its virtual servers unless the virtual server overrides the monitor selection. Likewise, the monitor defined for the pool is used to monitor all of its pool members, unless the pool member overrides the monitor selection.

It is important not to over-monitor a pool member. If a monitor is assigned to the server and/or virtual server and also to the pool and/or pool member, then both of the monitors fire, effectively monitoring the virtual server twice. In most cases, monitors should be configured at the server/virtual server or at the pool/pool member, but not both.

Prober pools

Prober pools allow the specification of particular set of BIG-IP devices that BIG-IP DNS may use to monitor a resource. This might be necessary in situations where a firewall is between certain BIG-IP systems and monitored resource, but not between other BIG-IP systems and those same resources. In this case, a prober pool can be configured and assigned to the server to limit probe requests to those BIG-IP systems that can reach the monitored resource. For more information about prober pools, refer to **About Prober** pools in ***BIG-IP Global Traffic Manager: Concepts***.

Note*: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Wide IPs

A wide IP maps a fully qualified domain name (FQDN) to one or more pools. The pools contain virtual servers. When a Local Domain Name Server (LDNS) makes a request for a domain that matches a wide IP, the configuration of the wide IP determines which virtual server address should be returned.

Wide IP names can contain the wildcard characters * (to match one or more characters) and ? (to match one character).

For more information about wide IPs, refer to **Wide IPs** in ***BIG-IP Global Traffic Manager: Concepts***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Record types

BIG-IP DNS 12.0 and later supports the following wide IP resource record types:

- A - Address Record
- AAAA - IPv6 Address Record
- CNAME - Canonical Name Record
- MX – Mail Exchanger
- SRV– Available Services
- NAPTR– Naming Authority Pointer Record

wide IP definitions consist of a Fully Qualified Domain Name (FQDN) and a Resource Record type. Each wide IP must have a RR type associated with it, and any pools attached to the wide IP must contain members of the same RR type.

***Note**: In BIG-IP 12.0 and later, pools can no longer be a mix of RR types. In BIG-IP 11.x through- 11.6, pools may contain both A and AAAA records.*

The notable exception to this are CNAME pools. A CNAME pool can be attached to any wide IP type, but it must either contain a static CNAME definition (FQDN) or a wide IP with the RR type that is the same as wide IP that has the CNAME pool attached.

The following figure shows a wide IP with a CNAME configured on the pool.

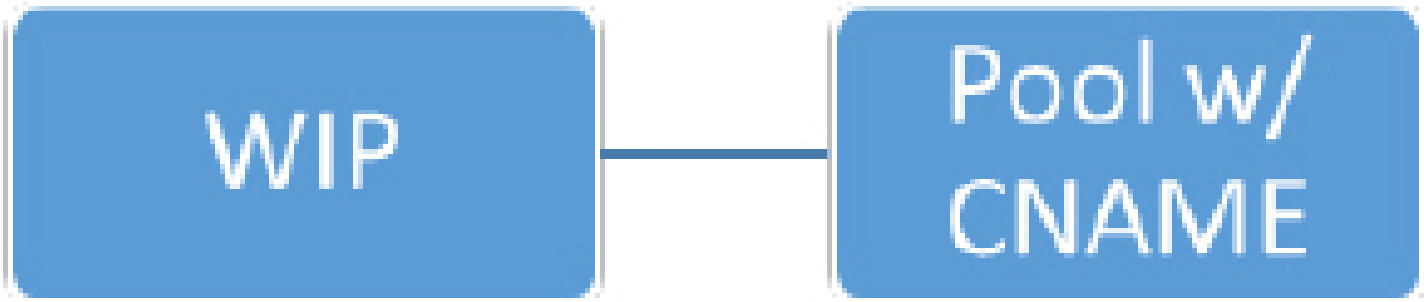

Figure 6.6 wide IP with CNAME pool configuration

Additional RR types are supported in versions earlier than BIG-IP 12.0, but they are configured in Zonerunner and are not wide IPs.

For more information on Zonerunner RR types, refer to **Using ZoneRunner to Configure DNS Zones** in ***BIG-IP DNS Services: Implementations***.

***Notes**:*

- *DNSSEC uses additional resource record types. However, these record types do not correspond to a specific wide IP configuration on the BIG-IP device.*
- ***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Wide IPs and pools

This section covers how to configure wide IPs and pools on BIG-IP devices in BIG-IP DNS 12.0 and later.

Pool definitions have changed significantly for the new resource record types. Pools attached to wide IPs can now contain other wide IP names as members. This change is necessary because resource records of types MX, SRV, NAPTR and CNAME return FQDNs rather than IP addresses.

wide IPs must be configured for each RR type it needs to answer. This means multiple wide IPs exist with the same hostname but different resource record types.

A and AAAA records

The following figures show how to configure resource records of types A and AAAA. In BIG-IP 12.0 and later, each RR type has its own wide IP to handle queries for that RR type.

Figure 6.7 “A” Resource Record configuration

Figure 6.8 “AAAA” Resource Record configuration

CNAME records

The following figures show how CNAMEs are configured for a type A resource record and for a static target.

The following figure shows a wide IP of type A configured. The wide IP has a pool of type CNAME associated with it. The CNAME pool has a single pool member that is not an IP address. Instead, the pool contains the name of another wide IP of type A. The CNAME pool must contain either another CNAME or a wide IP of the same type as the original wide IP requested.

Figure 6.9 CNAME For Type A Resource Record

The following figure shows a CNAME pool attached to a type A wide IP. In this instance, the CNAME has a Static Target defined. A Static Target can point to a FQDN not configured on the BIG-IP system. No health checks are applied to the Static Target.

Figure 6.10 CNAME For Type A Resource Record with Static Target

MX records

The following figure shows how wide IPs of type MX are configured.

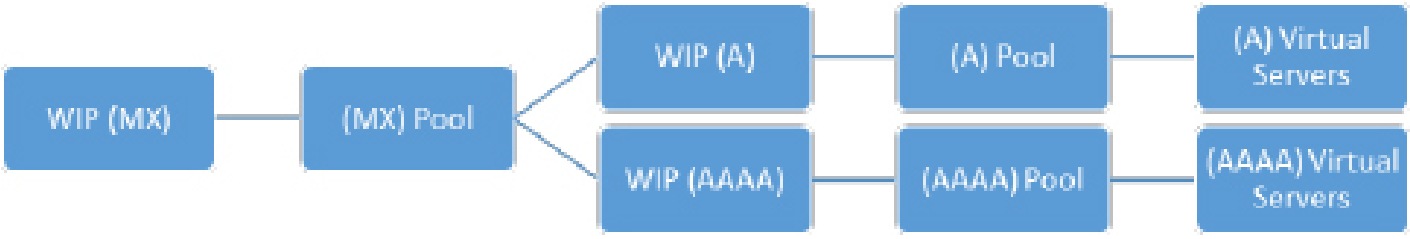

Figure 6.11 Type MX wide IP

In the first two boxes, the wide IP type is MX and the pool type is MX. The MX pool contains hostnames for other wide IPs configured on the BIG-IP.

Since type A and type AAAA records are valid MX hostnames, a single hostname in the MX pool points to any A or AAAA wide IPs configured with that hostname.

SRV records

The following figure shows how wide IPs of type SRV are configured. The left-most wide IP and pool attached to it both have type SRV. The SRV pool points to a wide IP also configured on the BIG-IP system.

Figure 6.12 SRV wide IPs configuration

NAPTR records

The following figure shows how wide IPs of type NAPTR are configured. The left-most wide IP and pool attached to it both have type NAPTR. A NAPTR pool can point to wide IPs on the same BIG-IP with types SRV, A or AAAA. In this example, the NAPTR pool contains both type SRV and type A wide IPs.

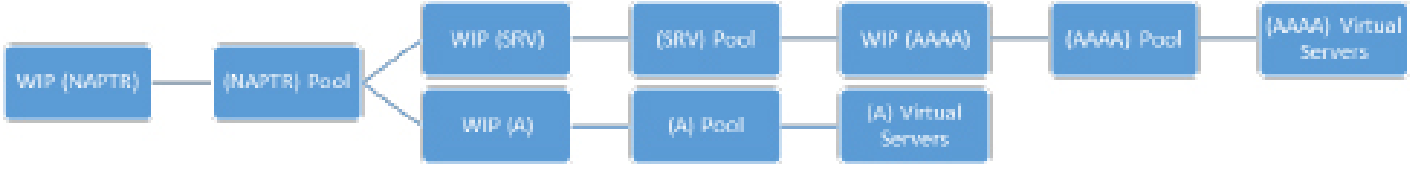

Figure 6.13 NAPTR type wide IP configuration

BIG-IP DNS Minimal Response Setting

In BIG-IP 12.0 and later, a configuration setting at the wide IP level called Minimal Response. This value determines whether DNS attempts resolve and return the IP addresses associated with hostnames contained in wide IP pools. The default value for Minimal Response is Enabled. F5 chose this default because it preserves the default behavior from BIG-IP versions prior to BIG-IP 12.0.

Examples

The following examples show the results of a **dig** command for the hostname _sip._udp.example.com. When Minimal Response is enabled, only the Answer Section is returned and contains the wide IP names contained in the SRV pool.\ The following example shows wide IP with Minimal Response enabled (Default)

Question Section

```
_sip._udp.example.com IN SRV
```

Answer Section:

```
_sip._udp.example.com.3222 IN SRV 20 10 6050 sip01.example.com

_sip._udp.example.com.3222 IN SRV 10 10 6050 sip02.example.com
```

When **Minimal Response** is disabled, the DNS resolves the hostnames in the **Answer** section and provides the information in the **Additional** section of the query response. The following example shows wide IP with **Minimal Response** disabled.

Question Section

```
_sip._udp.example.com IN SRV
```

Answer Section:

```
_sip._udp.example.com.3222 IN SRV 20 10 6050 sip01.example.com

_sip._udp.example.com.3222 IN SRV 10 10 6050 sip02.example.com
```

Additional Section

```
_sip._udp.example.com.3222 IN A 10.10.10.50

_sip._udp.example.com.3222 IN A 10.10.10.51
```

Effects on DNS Performance

A single query to a wide IP may include multiple additional name resolutions to wide IP pool members. Refer to NAPTR records for an example of a name resolution that requires multiple additional name resolutions.

Configuring the DNS with Minimal Response disabled can have performance impacts due to the additional steps the DNS must perform for each query to a wide IP.

F5 recommends setting **Minimal Response** to **Enabled**.

Load balancing logic

BIG-IP DNS provides a tiered load balancing mechanism for wide IP resolution. At the first tier, BIG-IP DNS chooses an appropriate pool of servers, and then, at the second tier, it chooses an appropriate virtual server.

For complete information about BIG-IP DNS load balancing, refer to **About load balancing and Global Traffic Manager** in ***BIG-IP Global Traffic Manager: Concepts***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

The wide IP has four static load balancing methods available for choice of an available pool.

The wide IP pool has several dynamic and static load balancing methods available to choose an available pool member.

The dynamic load balancing methods rely on metrics gathered to make a load balancing decision.

The static load balancing methods make a load balancing decision based on a set pattern. The pool allows the specification of preferred, alternate, and fallback load balancing options.

Not every load balancing method is available for each of the options.

When choosing an available pool member, the preferred method is tried first. When using a dynamic load balancing method as the preferred load balancing method, it is possible for load balancing to fail. For example, if the round-trip time method is chosen, when the first request arrives from an LDNS, there are no metrics available for it. Since there are no metrics available, the preferred method fails and BIG-IP DNS schedules the gathering of round-trip time metrics for that LDNS.

When the preferred method fails, the system falls back to the alternate method; therefore, when using a dynamic preferred method, it is important to specify an alternate method.

The fallback method is used to ensure that a resource is returned from the pool. The fallback method ignores the availability status of the resource being returned.

There are load balancing methods available that do not actually load balance. Methods such as none, drop packet, return to DNS, and fallback IP control the behavior of load balancing, but do not actually use the configured pool members.

Maximum number of available pool members

The BIG-IP DNS system can select multiple pool members in making load balancing decisions, and can use multiple pool members in a DNS response. In versions earlier than BIG-IP DNS 13.0.0, the maximum number of available pool members that the system can return in a response is limited to 16.

Beginning in BIG-IP DNS 13.0.0, the maximum number of available pool members that the system can return is 500.

The following GSLB pool types are affected by this change:

- A
- AAAA
- MX
- SRV
- NAPTR

You set the value for the wide IP in the **Maximum Answers Returned** attribute. Large DNS responses necessitate the switch to TCP and could affect performance.

***Note**: The number of pool members returned in a DNS response is the lower number between the maximum number chosen for the pool and the number of available pool members.*

Topology

BIG-IP DNS can make load balancing decisions based upon the geographical location of the LDNS making the DNS request. The location of the LDNS is determined from a GeoIP database. In order to use topology, the administrator must configure topology records describing how BIG-IP DNS should make its load balancing decisions. For more information on topology load balancing, refer to **Using Topology Load Balancing to Distribute DNS Requests to Specific Resources** in ***BIG-IP Global Traffic Manager: Load Balancing*** or K13412: Overview of BIG-IP DNS Topology records (11.x - 12.x).

Topology load balancing can be used to direct users to the servers that are geographically close, or perhaps to direct users to servers that have localized content.

BIG-IP system software provides a pre-populated database that provides a mapping of IP addresses to geographic locations. The administrator can also create custom group call regions. For example, it is possible to create a custom region that groups certain IP addresses together or that groups certain countries together.

Updates for the GeoIP database are provided on a regular basis. K11176: Downloading and installing updates to the IP geolocation database contains information about updating the database.

Topology records are used to map an LDNS address to a resource. The topology record contains three elements:

- A request source statement that specifies the origin LDNS of a DNS request.
- A destination statement that specifies the pool or pool member to which the weight of the topology record is assigned.
- A weight that the BIG-IP system assigns to a pool or a pool member during the load balancing process.

When determining how to load balance a request, BIG-IP DNS uses the object that has the highest weight according the matching topology records.

Important: When configuring topology load balancing at the wide IP level, topology records with a pool destination statement must exist. Other destination statement types (such as data center or country) may be used when using topology as a pool level load balancing method.

BIG-IP DNS Architectures

This section describes three common deployment methodologies for using a BIG-IP system in a DNS environment. For more information, refer to ***BIG-IP Global Traffic Manager: Implementations***.

Note*: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Delegated mode

When operating in delegated mode, requests for wide IP resource records are redirected or delegated to BIG-IP DNS. The BIG-IP system does not see all DNS requests, and operates on requests for records that are sent to it.

For more information, refer to ***BIG-IP Global Traffic Manager: Implementations***.

Note*: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Screening mode

When operating in screening mode, BIG-IP DNS sits in front of one or more DNS servers. This configuration allows for easy implementation of additional BIG-IP system features for DNS traffic because DNS requests for records other than wide IPs pass through BIG-IP DNS. If the request matches a wide IP, BIG-IP DNS responds to

the request. Otherwise, the request is forwarded to the DNS servers. This configuration can provide the following benefits:

- DNS query validation: When a request arrives at BIG-IP DNS, BIG-IP DNS validates that the query is well formed. BIG-IP DNS can drop malformed queries, protecting the back-end DNS servers from seeing the malformed queries.
- DNSSEC dynamic signing: When the responses from the DNS server pass back through BIG-IP DNS, it is possible for BIG-IP DNS to sign the response. This allows the use of DNSSEC with an existing zone and DNS servers, but takes advantage of any cryptographic accelerators in a BIG-IP device.
- Transparent Caching: When the responses from the DNS server pass back through the BIG-IP system , it can cache the response. Future requests for the same records can be served directly from the BIG-IP system reducing the load on the back-end DNS servers.

For more information, refer to ***BIG-IP Global Traffic Manager: Implementations***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Replacing a DNS server

It is also possible for the BIG-IP system to operate as a stand-alone, authoritative DNS server for one or more zones. In this configuration, all DNS requests for a zone are sent to BIG-IP DNS. Any requests for a wide IP are handled by BIG-IP DNS and other requests are sent to the local bind instance on the BIG-IP system. ZoneRunner is used to manage the records in the local-bind instance. For more information, refer to **Replacing a DNS Server with BIG-IP DNS** in ***BIG-IP Global Traffic Manager: Implementations***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

BIG-IP DNS iQuery

- iQuery is an XML protocol that BIG-IP systems use to communicate with each other. BIG-IP DNS uses iQuery for various tasks:
- Determining the health of objects in BIG-IP DNS configuration.
- Exchanging information about BIG-IP DNS synchronization group state.
- Providing a transport for synchronizing BIG-IP DNS configuration throughout the synchronization group.
- Communicating LDNS path probing metrics.
- Exchanging wide IP persistence information.
- Gathering BIG-IP system configuration when using auto-discovery.

All of these tasks combined provide a BIG-IP DNS synchronization group with a unified view of BIG-IP DNS configuration and state.

iQuery Agents

All BIG-IP DNS devices are iQuery clients. The **gtmd** process on each BIG-IP DNS device connects to the **big3d** process on every BIG-IP server defined in BIG-IP DNS configuration, which includes both BIG-IP DNS and BIG-IP LTM.

These are long-lived connections made using TCP port 4353. This set of connections among BIG-IP DNS devices and between BIG-IP DNS and BIG-IP LTM devices is called an iQuery mesh.

iQuery communication is encrypted using SSL. The devices involved in the communication authenticate each other using SSL certificate-based authentication. For information, refer to **Communications Between BIG-IP DNS and Other Systems** in ***BIG-IP Global Traffic Manager: Concepts***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

To monitor the health of objects in the configuration, BIG-IP DNS devices in the synchronization group sends monitor requests using iQuery to another iQuery server that is closer to the target of the monitor. All BIG-IP DNS devices in the synchronization group agrees on which BIG-IP DNS is responsible for initiating the monitoring request. The result of the monitoring request is sent by the iQuery server to all BIG-IP DNS devices connected to it that are participating in the synchronization group.

***Note**: A lack of a unified view of the iQuery mesh causes unpredictable behavior. For example, if each BIG-IP DNS device is not connected to the same set of other BIG-IP DNS devices, there can be disagreement of monitor responsibility resulting in object availability flapping ("flapping" is when a device is marked down and up repeatedly).*

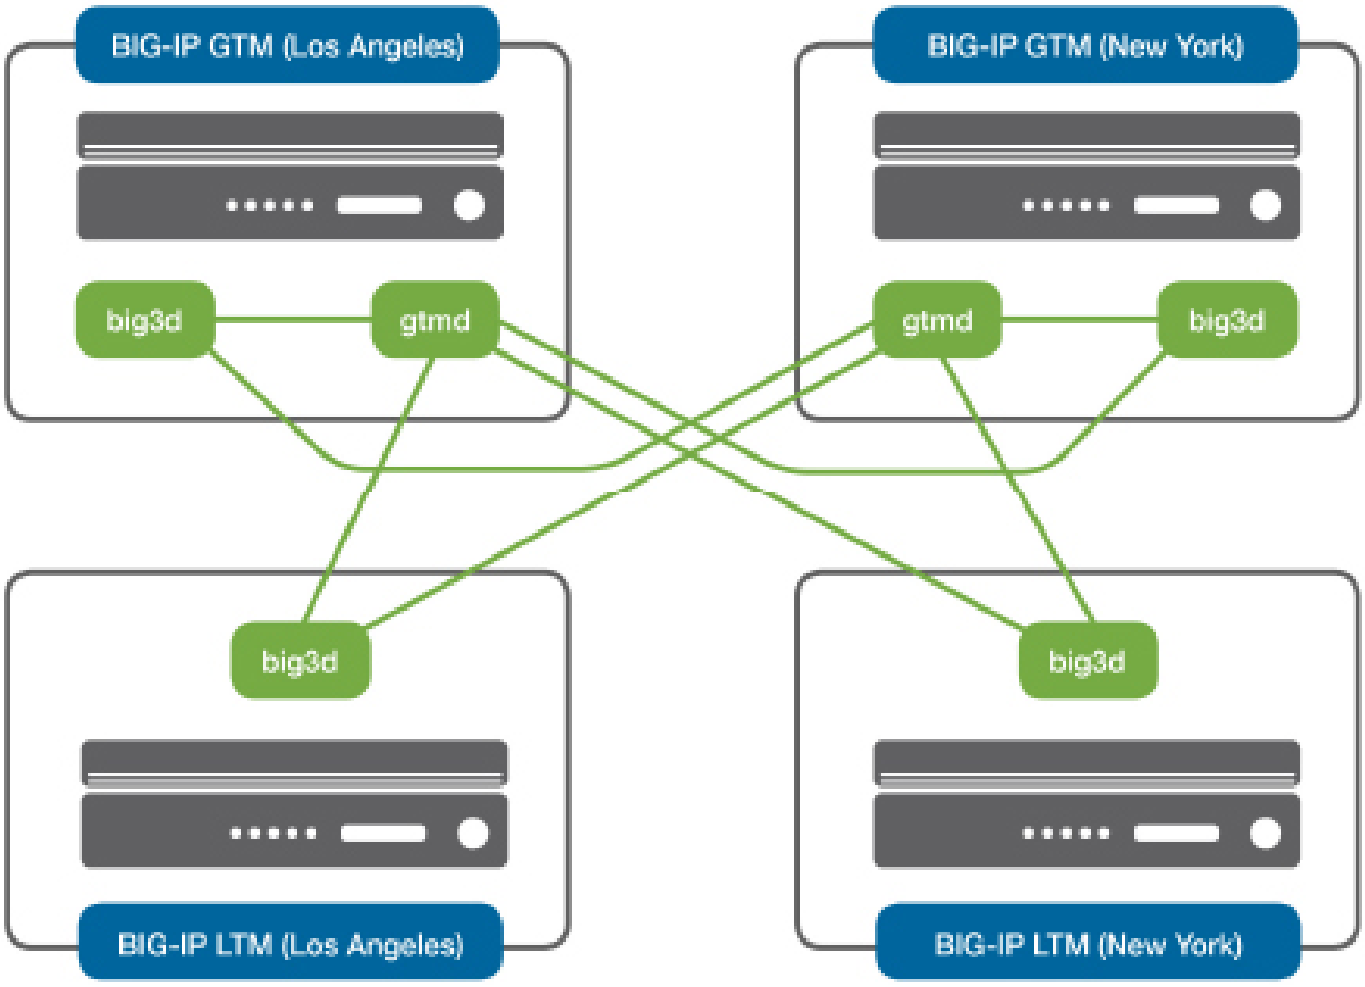

Figure 6.14 BIG-IP LTM/BIG-IP DNS configuration

big3d software versioning

The version of the **big3d** software installed on each device must be the same or later than the version of software used on BIG-IP DNS devices. For information on updating the **big3d** software, refer to **Running the big3d_install script** in ***BIG-IP Global Traffic Manager: Implementations***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

Determining the health of the iQuery mesh

Reviewing log files or SNMP traps

The **/var/log/gtm** log file contains information about connection status changes to big3d agents. When a new connection is established to a **big3d** agent or when a connection is lost, a log message is generated.

Example of connection-lost messages:

```
alert gtmd[8663]: 011a500c:1: SNMP_TRAP: Box 10.14.20.209 state change green --> red (Box 10.14.20.209 on Unavailable)
alert gtmd[8663]: 011a5004:1: SNMP_TRAP: Server /Common/gtm-3 (ip=10.14.20.209) state change green --> red (No communication)
```

Example of connection-established messages:

```
alert gtmd[8663]: 011a500b:1: SNMP_TRAP: Box 10.14.20.209 state change red --> green
alert gtmd[8663]: 011a5003:1: SNMP_TRAP: Server /Common/gtm-3 (ip=10.14.20.209) state change red --> green
```

If a connection to a configured BIG-IP server is down, repeated **Connection in progress to** messages are generated:

For example:

```
notice gtmd[8663]: 011ae020:5: Connection in progress to 10.14.20.209
notice gtmd[8663]: 011ae020:5: Connection in progress to 10.14.20.209
notice gtmd[8663]: 011ae020:5: Connection in progress to 10.14.20.209
```

tmsh show gtm iquery command

You can use **tmsh show gtm iquery** command to display that status of all of the iQuery connections on a BIG-IP DNS device. The command displays each IP address:

```
# tmsh show gtm iquery
--------------------------------------
Gtm::IQuery: 10.12.20.207
--------------------------------------
Server gtm-1
Data Center DC1
iQuery State connected
Query Reconnects 1
Bits In 8.2M
Bits Out 47.7K
Backlogs 0
Bytes Dropped 96
Cert Expiration Date 10/29/24 04:38:53
Configuration Time 12/08/14 16:37:49
--------------------------------------
Gtm::IQuery: 10.14.20.209
--------------------------------------
```

```
Server gtm-3
Data Center DC2
iQuery State connected
Query Reconnects 0
Bits In 8.2M
Bits Out 45.7K
Backlogs 0
Bytes Dropped 0
Cert Expiration Date 10/29/24 04:38:53
Configuration Time 12/08/14 16:37:49
```

For more information, refer to K13690: Troubleshooting BIG-IP DNS synchronization and iQuery connections (11.x - 15.x).

iqdump

You can use the **iqdump** command to check the communication path and SSL certificate-based authentication from a BIG-IP DNS to another device in the iQuery mesh.

The syntax of the **iqdump** command is **iqdump <ip address> <synchronization group name>**. When using the **iqdump** command, BIG-IP DNS synchronization group name is optional.

For example:

```
# iqdump 10.14.20.209
<!-- Local hostname: gtm1.example.com -->
<!-- Connected to big3d at::ffff:10.14.20.209:4353 -->
<!-- Subscribing to syncgroup: default -->
<!-- Mon Dec 8 16:45:00 2014 -->
<xml_connection>
<version>11.4.0</version>
<big3d>big3d Version 11.5.1.6.0.312</big3d>
<kernel>linux</kernel>
...
```

***Note**: The **iqdump** command continues to run until it is interrupted by pressing Ctrl-C.*

If there is a problem with the communication path or the SSL authentication, the **iqdump** command fails and reports an error.

The version of BIG-IP software being run on the remote system is reported in the version XML stanza. The version of the **big3d** software running on the remote system is reported in the **<big3d>** XML stanza.

For more information, refer to K13690: Troubleshooting BIG-IP DNS synchronization and iQuery connections (11.x - 15.x).

BIG-IP DNS device service clustering

In BIG-IP 13.0 and later, BIG-IP DNS includes full support for device service clustering (DSC). In previous versions, devices configured in a DSC failover device group are not fully supported when more than two BIG-IP systems are included in the cluster.

For example, if a server object is configured in a BIG-IP redundant configuration, virtual server auto-discovery does not function properly but instead adds then deletes non-floating virtual services in your device group.

In BIG-IP 13.0 and later, BIG-IP DNS creates a server object representing any number of BIG-IP systems in the cluster. Virtual servers are tracked both by the server and the device; therefore, a virtual server is removed from your device group only if it is first removed from all devices which interacted with it.

BIG-IP DNS query logging

BIG-IP DNS can log debugging information about the decision-making process when resolving a wide IP. These logs can report which pool and virtual server were chosen for a wide IP resolution or why BIG-IP DNS is unable to respond.

For information about enabling query logging, refer to K14615: Configuring the BIG-IP DNS system to log wide IP request information .

***Note**: Query logging should only be enabled only for troubleshooting and not for long periods of time*.

For DNS configurations that provide services beyond wide IP resolution, for example DNS recursive resolvers or DNS Express, it is possible to enable DNS query and response logging. For more information, refer to **External Monitoring** in ***BIG-IP Systems: Implementations***.

***Note**: For information about how to locate F5 product manuals, refer to K98133564: Tips for searching AskF5 and finding product documentation.*

BIG-IP DNS Statistics

The BIG-IP system maintains statistics for objects throughout the system. Due to the variety of statistics gathered and the breadth of configuration elements covered, it is impractical to cover them within this manual. Statistics are documented throughout the user manuals where they are featured.

You can view statistics in the Configuration utility by going to **Statistics** > **Module Statistics**. In **tmsh**, statistics are visible using the show command with particular objects. Typically, statistics are either gauge or counters. A gauge keeps track of a current state, for example current connections. A counter keeps an incrementing count of how many times a particular action is taken, for example total requests.

Reviewing statistics that use counters may provide insight into the proportion of traffic which is valid, or perhaps may indicate that there is a configuration error.

For example, comparing the **Dropped Queries** counter to **Total Queries** counter shows that there are some drops, but this may not be of concern because the drops are fairly low as a percentage of the total.

```
# show gtm wide IP www.example.com
---------------------------
Gtm::wide IP: www.example.com
---------------------------
Status
Availability : available
State : enabled
Reason : Available
Requests
Total 1765
```

```
A 1552
AAAA 213
Persisted 0
Resolved 1762
Dropped 3
Load Balancing
Preferred 1760
Alternate 2
Fallback 0
CNAME Resolutions 0
Returned from DNS 0
Returned to DNS 0
```

Statistics are also available for polling using SNMP and can be polled, cataloged over time, and graphed by a Network Management Station (NMS).

BIG-IP LTM-DNS operations guide

- Chapter 1: Guide introduction and contents
- Chapter 2: BIG-IP LTM Load Balancing
- Chapter 3: BIG-IP LTM network address objects
- Chapter 4: BIG-IP LTM Virtual Servers
- Chapter 5: BIG-IP LTM Profiles
- Chapter 7: iRules
- Chapter 8: Logging

Supplemental Information

- About operations guides
- Optimizing the support experience

Applies to:

Product: BIG-IP, BIG-IP DNS, BIG-IP LTM
15.X.X, 14.X.X, 13.X.X, 12.X.X, 11.6.X, 11.5.X